

EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	2553	707/2.ccls.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/09/27 13:56
L2	6	top-k.clm.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/09/27 14:00
L3	0	1 and 2	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/09/27 14:00
L4	7296	707/3.ccls.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/09/27 14:00
L5	1	4 and 2	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/09/27 14:03
L6	1103	707/7.ccls.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/09/27 14:03
L7	3	6 and 2	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/09/27 14:06

EAST Search History

L10	0	7 and (top-k near join near quer\$3)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/09/27 14:07
L11	1	4 and (top-k near join near quer\$3)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/09/27 14:07


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☐ The ACM Digital Library ☒ The Guide

TOP-K QUERY

SEARCH

THE GUIDE TO COMPUTING LITERATURE

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)
Terms used **TOP K QUERY**

Found 96 of 955,089

Sort results by

relevance ☒

Display results

expanded form ☒

Save results to a Binder

Search Tips

☐ Open results in a new window

Try an Advanced Search

Try this search in [The Digital Library](#)

Results 1 - 20 of 96

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [next](#)Relevance scale ☐ ☐ ☐ ☐ ☐
1 [Reducing network traffic in unstructured P2P systems using Top-k queries](#)

Reza Akbarinia, Esther Pacitti, Patrick Valduriez

May 2006 **Distributed and Parallel Databases**, Volume 19 Issue 2-3

Publisher: Kluwer Academic Publishers

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

A major problem of unstructured P2P systems is their heavy network traffic. This is caused mainly by high numbers of query answers, many of which are irrelevant for users. One solution to this problem is to use Top-k queries whereby the user can specify a limited number (k) of the most relevant answers. In this paper, we present FD, a (Fully Distributed) framework for executing Top-k queries ...

Keywords: Network traffic, Peer-to-peer, Top-k query, Unstructured

2 [Data streams: Continuous monitoring of top-k queries over sliding windows](#)


Kyriakos Mouratidis, Spiridon Bakiras, Dimitris Papadias

June 2006 **Proceedings of the 2006 ACM SIGMOD international conference on Management of data SIGMOD '06**

Publisher: ACM Press

Full text available: pdf(576.03 KB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

Given a dataset P and a preference function f , a top-k query retrieves the k tuples in P with the highest scores according to f . Even though the problem is well-studied in conventional databases, the existing methods are inapplicable to highly dynamic environments involving numerous long-running queries. This paper studies continuous monitoring of top-k queries over a fixed-size window W of the most recent data. The window size can be expressed either in ...

Keywords: continuous queries, sliding windows, top-k processing

3 [Research session: DB and IR #1: KLEE: a framework for distributed top-k query algorithms](#)

Sebastian Michel, Peter Triantafillou, Gerhard Weikum

August 2005 **Proceedings of the 31st international conference on Very large data bases VLDB '05**

Publisher: VLDB Endowment

Full text available:  pdf(422.08 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper addresses the efficient processing of top-k queries in wide-area distributed data repositories where the index lists for the attribute values (or text terms) of a query are distributed across a number of data peers and the computational costs include network latency, bandwidth consumption, and local peer work. We present KLEE, a novel algorithmic framework for distributed top-k queries, designed for high performance and flexibility. KLEE makes a strong case for approximate top-k algor ...


4 [Research sessions: query processing II: Minimal probing: supporting expensive predicates for top-k queries](#)



Kevin Chen-Chuan Chang, Seung-won Hwang

June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data SIGMOD '02**

Publisher: ACM Press

Full text available:  pdf(1.53 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper addresses the problem of evaluating ranked *top-k* queries with expensive predicates. As major DBMSs now all support expensive user-defined predicates for Boolean queries, we believe such support for ranked queries will be even more important: First ranked queries often need to model user-specific concepts of preference, relevance, or similarity, which call for dynamic user-defined functions. Second, middleware systems must incorporate external predicates for integrating autonomo ...

5 [P2P and network algorithms: Efficient top-K query calculation in distributed networks](#)



Pei Cao, Zhe Wang

July 2004 **Proceedings of the twenty-third annual ACM symposium on Principles of distributed computing**

Publisher: ACM Press

Full text available:  pdf(206.90 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper presents a new algorithm to answer top-k queries (e.g. "find the *k* objects with the highest aggregate values") in a distributed network. Existing algorithms such as the Threshold Algorithm [10] consume an excessive amount of bandwidth when the number of nodes, *m*, is high. We propose a new algorithm called "Three-Phase Uniform Threshold" (TPUT). TPUT reduces network bandwidth consumption by pruning away ineligible objects, and terminates in three round-trips regard ...

Keywords: distributed networks, instance optimality, top-k algorithms

6 [Efficiency: Efficient and self-tuning incremental query expansion for top-k query processing](#)



Martin Theobald, Ralf Schenkel, Gerhard Weikum

August 2005 **Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval SIGIR '05**

Publisher: ACM Press

Full text available:  pdf(532.92 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We present a novel approach for efficient and self-tuning query expansion that is embedded into a top-k query processor with candidate pruning. Traditional query expansion methods select expansion terms whose thematic similarity to the original query terms is above some specified threshold, thus generating a disjunctive query with much higher dimensionality. This poses three major problems: 1) the need for hand-tuning the expansion threshold, 2) the potential topic dilution with overly aggressiv ...

Keywords: incremental merge, probabilistic candidate pruning, query expansion, top-k

ranking

7 Evaluating top- k queries over web-accessible databases

Amélie Marian, Nicolas Bruno, Luis Gravano

June 2004 **ACM Transactions on Database Systems (TODS)**, Volume 29 Issue 2**Publisher:** ACM PressFull text available: pdf(1.03 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

A query to a web search engine usually consists of a list of keywords, to which the search engine responds with the best or "top" k pages for the query. This top- k query model is prevalent over multimedia collections in general, but also over plain relational data for certain applications. For example, consider a relation with information on available restaurants, including their location, price range for one diner, and overall food rating. A user who queries such a relation might ...

Keywords: Parallel query processing, query optimization, top- k query processing, web databases.

8 Top- k selection queries over relational databases: Mapping strategies and performance evaluation

Nicolas Bruno, Surajit Chaudhuri, Luis Gravano

June 2002 **ACM Transactions on Database Systems (TODS)**, Volume 27 Issue 2**Publisher:** ACM PressFull text available: pdf(1.64 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In many applications, users specify target values for certain attributes, without requiring exact matches to these values in return. Instead, the result to such queries is typically a rank of the "top k " tuples that best match the given attribute values. In this paper, we study the advantages and limitations of processing a top- k query by translating it into a single range query that a traditional relational database management system (RDBMS) can process efficiently. In particular, ...

Keywords: Multidimensional histograms, top- k query processing

9 In-network processing and routing: The threshold join algorithm for top- k queries in distributed sensor networks

D. Zeinalipour-Yazti, Z. Vagena, D. Gunopulos, V. Kalogeraki, V. Tsotras, M. Vlachos, N. Koudas, D. Srivastava

August 2005 **Proceedings of the 2nd international workshop on Data management for sensor networks DMSN '05****Publisher:** ACM PressFull text available: pdf(162.07 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In this paper we present the *Threshold Join Algorithm (TJA)*, which is an efficient TOP- k query processing algorithm for distributed sensor networks. The objective of a top- k query is to find the k highest ranked answers to a user defined similarity function. The evaluation of such a query in a sensor network environment is associated with the transfer of data over an extremely expensive communication medium. *TJA* uses a non-uniform threshold on the queried attribute in order ...

Keywords: distributed systems, sensor networks, top-K queries

10 Evaluating Top-k Queries over Web-Accessible Databases

Amelie Marian

February 2002 **Proceedings of the 18th International Conference on Data Engineering
ICDE '02****Publisher:** IEEE Computer SocietyFull text available:  [Publisher Site](#) Additional Information: [full citation](#), [abstract](#), [citations](#)

A query to a web search engine usually consists of a list of keywords, to which the search engine responds with the best or "top" k pages for the query. This top-k query model is prevalent over multimedia collections in general, but also over plain relational data for certain applications. For example, consider a relation with information on available restaurants, including their location, price range for one diner, and overall food rating. A user who queries such a relation might simply specif ...

11 A Sampling-Based Estimator for Top-k QueryFebruary 2002 **Proceedings of the 18th International Conference on Data Engineering
ICDE '02****Publisher:** IEEE Computer SocietyFull text available:  [Publisher Site](#) Additional Information: [full citation](#), [abstract](#), [citations](#)

Top-k queries arise naturally in many database applications that require searching for records whose attribute values are close to those specified in a query. In this paper, we study the problem of processing a top-k query by translating it into an approximate range query that can be efficiently processed by traditional relational DBMSs. We propose a sampling-based approach, along with various query mapping strategies, to determine a range query that yields high recall with low access cost. Our e ...

Keywords: Top-K query, Sampling, Range query**12 Adaptive Processing of Top-k Queries in XML**

Amelie Marian, Sihem Amer-Yahia, Nick Koudas, Divesh Srivastava

April 2005 **Proceedings of the 21st International Conference on Data Engineering
(ICDE'05) - Volume 00 ICDE '05****Publisher:** IEEE Computer SocietyFull text available:  [Publisher Site](#) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

The ability to compute top-k matches to XML queries is gaining importance due to the increasing number of large XML repositories. The efficiency of top-k query evaluation relies on using scores to prune irrelevant answers as early as possible in the evaluation process. In this context, evaluating the same query plan for all answers might be too rigid because, at any time in the evaluation, answers have gone through the same number and sequence of operations, which limits the speed at which score ...

13 Research sessions: Web, XML and IR: FlexPath: flexible structure and full-text querying for XML

Sihem Amer-Yahia, Laks V. S. Lakshmanan, Shashank Pandit

June 2004 **Proceedings of the 2004 ACM SIGMOD international conference on
Management of data****Publisher:** ACM PressFull text available:  [pdf\(437.86 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

Querying XML data is a well-explored topic with powerful database-style query languages such as XPath and XQuery set to become W3C standards. An equally compelling paradigm for querying XML documents is full-text search on textual content. In this paper, we study fundamental challenges that arise when we try to integrate these two querying

paradigms. While keyword search is based on approximate matching, XPath has exact match semantics. We address this mismatch by considering queries on structure ...

14 Demo session: advanced applications: RankSQL: supporting ranking queries in relational database management systems

Chengkai Li, Mohamed A. Soliman, Kevin Chen-Chuan Chang, Ihab F. Ilyas

August 2005 **Proceedings of the 31st international conference on Very large data bases VLDB '05**

Publisher: VLDB Endowment

Full text available:  pdf(201.20 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Ranking queries (or top-k queries) are dominant in many emerging applications, e.g., similarity queries in multimedia databases, searching Web databases, middleware, and data mining. The increasing importance of top-k queries warrants an efficient support of ranking in the relational database management system (RDBMS) and has recently gained the attention of the research community. Top-k queries aim at providing only the top k query results, according to ...

15 Research session: DB and IR #1: An efficient and versatile query engine for TopX search

Martin Theobald, Ralf Schenkel, Gerhard Weikum

August 2005 **Proceedings of the 31st international conference on Very large data bases VLDB '05**

Publisher: VLDB Endowment

Full text available:  pdf(442.21 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)


This paper presents a novel engine, coined *TopX*, for efficient ranked retrieval of XML documents over semistructured but nonschematic data collections. The algorithm follows the paradigm of threshold algorithms for top-k query processing with a focus on inexpensive sequential accesses to index lists and only a few judiciously scheduled random accesses. The difficulties in applying the existing top-k algorithms to XML data lie in 1) the need to consider scores for XML elements while aggregating ...

16 A Query Processing Mechanism for Top-k Query in P2P Networks

Hidekazu MATSUNAMI, Tsutomu TERADA, Shojiro NISHIO

April 2005 **Proceedings of the 21st International Conference on Data Engineering Workshops ICDEW '05**

Publisher: IEEE Computer Society

Full text available:  [Publisher Site](#) Additional Information: [full citation](#), [abstract](#)

Recently, there has been an increasing interest in content sharing on peer-to-peer (P2P) networks. Since such a system employs a flooding mechanism for queries and because each peer returns many search results, the system's response to a query creates heavy traffic. Therefore, we propose a new and more efficient query processing method for top-k queries on P2P networks. We focus on the fact that users usually need search results only with a higher score. Our method reduces the reply traffic by ...

17 Evaluation of top-k queries over structured and semi-structured data

Amelie Marian, Luis Gravano

January 2005 Doctoral Thesis

Publisher: Columbia University

Additional Information: [full citation](#), [abstract](#), [index terms](#)

This thesis addresses fundamental issues in defining and efficiently processing top-k queries for a variety of scenarios, presenting different query processing challenges. In all these scenarios, our query processing algorithms attempt to focus on the objects that are


most likely to be among the top-k matches for a given query, and discard---as early as possible---objects that are guaranteed not to qualify for the top- k answer, thus minimizing ...

18 Monitoring Top-k Query in Wireless Sensor Networks

Minji Wu, Jianliang Xu, Xueyan Tang, Wang-Chien Lee

April 2006 **Proceedings of the 22nd International Conference on Data Engineering (ICDE'06) - Volume 00 ICDE '06**

Publisher: IEEE Computer Society

Full text available:  [Publisher Site](#) Additional Information: [full citation](#), [abstract](#)

Top-k monitoring is important to many wireless sensor applications. This paper exploits the semantics of top-k query and proposes a novel energy-efficient monitoring approach, called FILA. The basic idea is to install a filter at each sensor node to suppress unnecessary sensor updates. The correctness of the top-k result is ensured if all sensor nodes perform updates according to their filters. We show via simulation that FILA outperforms the existing TAGbased approach by an order of magnitude.

19 A Sampling-Based Approach to Optimizing Top-k Queries in Sensor Networks

Adam Silberstein Silberstein, Rebecca Braynard, Carla Ellis, Kamesh Munagala, Jun Yang

April 2006 **Proceedings of the 22nd International Conference on Data Engineering (ICDE'06) - Volume 00 ICDE '06**

Publisher: IEEE Computer Society

Full text available:  [Publisher Site](#) Additional Information: [full citation](#), [abstract](#)

Wireless sensor networks generate a vast amount of data. This data, however, must be sparingly extracted to conserve energy, usually the most precious resource in battery-powered sensors. When approximation is acceptable, a model-driven approach to query processing is effective in saving energy by avoiding contacting nodes whose values can be predicted or are unlikely to be in the result set. To optimize queries such as top-k, however, reasoning directly with models of joint probability distribu ...

20 Research papers: optimization: RankSQL: query algebra and optimization for relational top-k queries

Chengkai Li, Kevin Chen-Chuan Chang, Ihab F. Ilyas, Sumin Song

June 2005 **Proceedings of the 2005 ACM SIGMOD international conference on Management of data**

Publisher: ACM Press

Full text available:  [pdf\(741.54 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

This paper introduces RankSQL, a system that provides a systematic and principled framework to support efficient evaluations of ranking (*top-k*) queries in relational database systems (RDBMS), by extending relational algebra and query optimization. Previously, *top-k* query processing is studied in the middleware scenario or in RDBMS in a "piecemeal" fashion, *i.e.*, focusing on specific operator or sitting outside the core of query engines. In contrast, we aim to support ranking ...

Results 1 - 20 of 96

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

[Sign in](#)[Web](#) [Images](#) [Video](#) ^{New!} [News](#) [Maps](#) [more »](#)

TOP-K QUERY

Search

[Advanced Search](#)
[Preferences](#)**Web**Results 1 - 10 of about 79,600 for **TOP-K QUERY**. (0.28 seconds)**[PDF] Efficient Top-K Query Calculation in Distributed Networks**File Format: PDF/Adobe Acrobat - [View as HTML](#)brary of **top-k** calculations for distributed comput-. ing infrastructures such as the PlanetLab.Lastly, we. plan to incorporate **top-k query** calculation ...crypto.stanford.edu/~cao/topk.pdf - [Similar pages](#)**Efficient Top-K Query Calculation in Distributed Networks**Efficient **Top-K Query** Calculation in Distributed Networks ... This paper presents a new algorithm to answer **top-k** queries (eg ``find the k objects with the ...crypto.stanford.edu/~cao/topk.html - 2k - [Cached](#) - [Similar pages](#)**[PPT] Top-k Query Processing**File Format: Microsoft Powerpoint - [View as HTML](#)**Top-k Query** Processing. Optimal aggregation algorithms for middleware. Ronald Fagin, Amnon Lotem, and Moni Naor. + Sushruth P. + Arjun Dasgupta ...

crystal.uta.edu/~gdas/Courses/Courses/

Spring2006/DBExploration/Arjun_Sushruth_fagin_ta.ppt - [Similar pages](#)**[PPT] RankSQL: Query Algebra and Optimization for Relational Top- k ...**File Format: Microsoft Powerpoint - [View as HTML](#)**Top-k** queries provides only the **top k query** results according to a user-specified ... **Top-k** queries are not treated as first class **query** type in RDBMS. ...

crystal.uta.edu/~gdas/Courses/Courses/Spring2006/DBExploration/RankSQL.ppt -

[Similar pages](#)**[PDF] Top- k Query Evaluation for Schema-Based Peer-to-Peer Networks**File Format: PDF/Adobe Acrobat - [View as HTML](#)success of ranking algorithms in Web search engine and **top-k query** evaluation.algorithms in databases, we propose a decentralized **top-k query** evaluation ...www.kbs.uni-hannover.de/Arbeiten/Publikationen/2004/topk_iswc.pdf - [Similar pages](#)**Top-k Query Evaluation for Schema-Based Peer-to-Peer Networks ...**Increasing the number of peers in a peer to peer network usually increases the number of answers to a given **query** as well. While having more answers is nice ...citeseer.ist.psu.edu/thaden04topk.html - 23k - [Cached](#) - [Similar pages](#)**[PDF] Monitoring Top- k Query in Wireless Sensor Networks**

File Format: PDF/Adobe Acrobat

A basic implementation of monitoring **top-k query**. would be to use a centralized approach where all sensor. readings are collected by the base station, ...ieeexplore.ieee.org/iel5/10757/33902/01617511.pdf - [Similar pages](#)**Takumi Okazaki: Retrieving subset of result before completing top ...**One approach to reducing **query** cost is to search only the **top-k** elements, ... The **top-k query** is a **query** which finds the k objects that have the highest ...dbpubs.stanford.edu/pub/2005-22 - 23k - [Cached](#) - [Similar pages](#)**[PPT] KLEE: A Framework for Distributed Top-k Query Algorithms**File Format: Microsoft Powerpoint - [View as HTML](#)

KLEE: A Framework for Distributed **Top-k Query** Algorithms. KLEE: Key Ideas. if mink / m is small TPUT retrieves a lot of data in Phase 2 ...
www.vldb2005.org/program/slides/thu/s637-michel.ppt - [Similar pages](#)

[PDF] **A Query Processing Mechanism for Top-k Query in P2P Networks**

File Format: PDF/Adobe Acrobat

Kalnis[5] proposes a system that realizes a **top-k query**. on a Gnutella-type P2P network. ...
ecuting a **top-k query**, each peer receiving the **query** returns ...
doi.ieeecomputersociety.org/10.1109/ICDE.2005.167 - [Similar pages](#)

Goooooooooooooogle ►

Result Page: 1 2 3 4 5 6 7 8 9 10 [Next](#)

Free! Speed up the web. [Download the Google Web Accelerator.](#)

TOP-K QUERY

Search

[Search within results](#) | [Language Tools](#) | [Search Tips](#) | [Dissatisfied? Help us improve](#)

[Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

©2006 Google